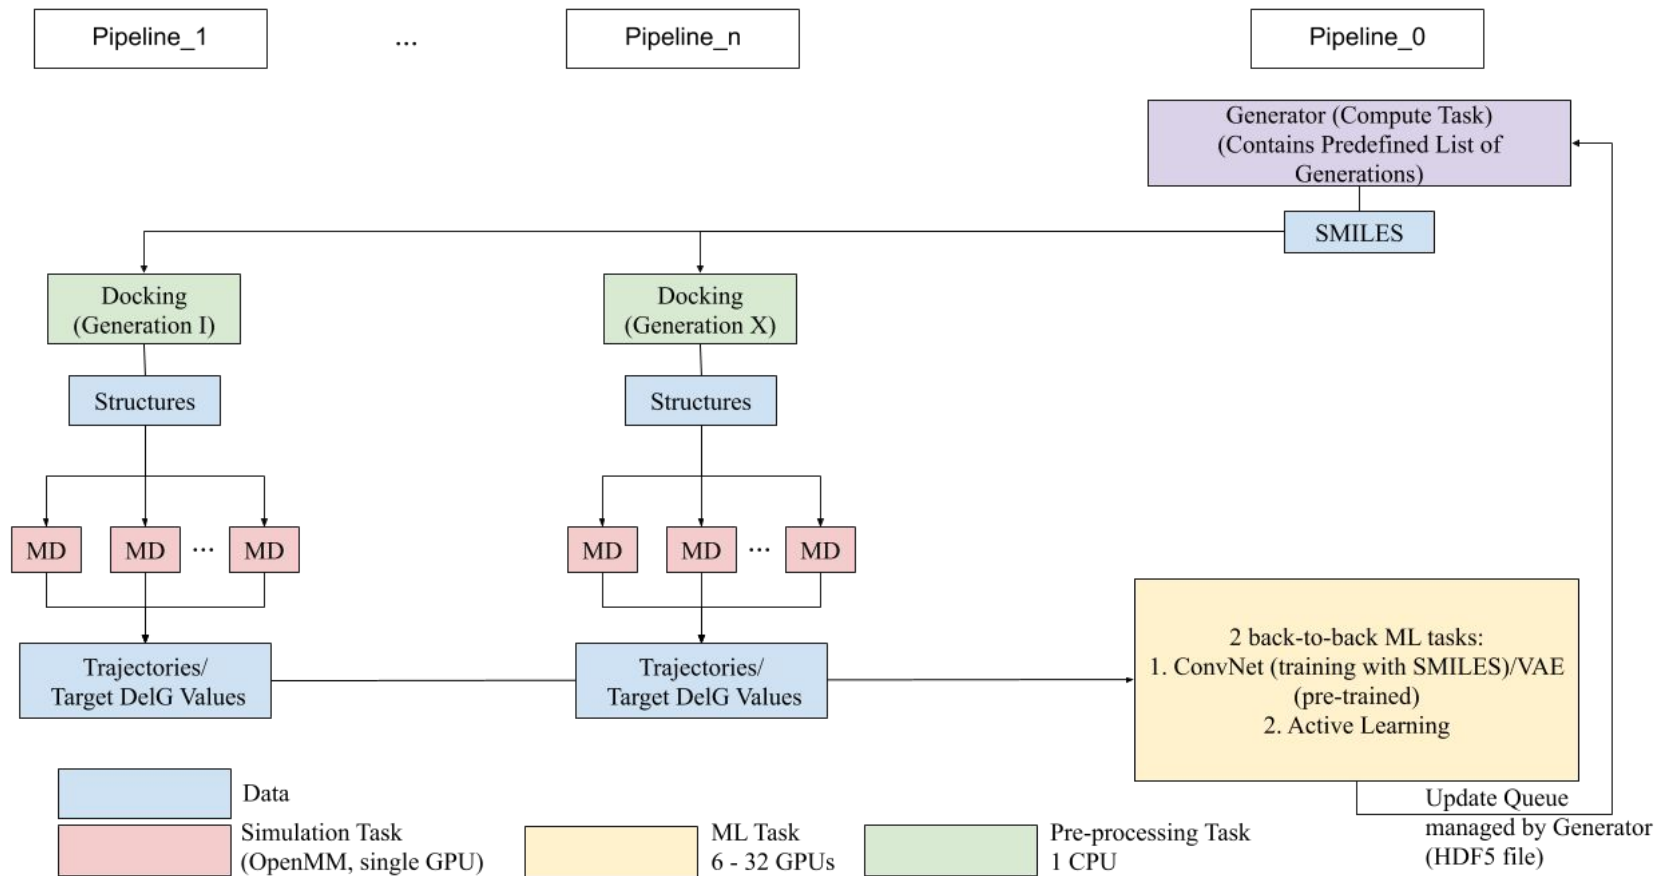


NGE Project Report

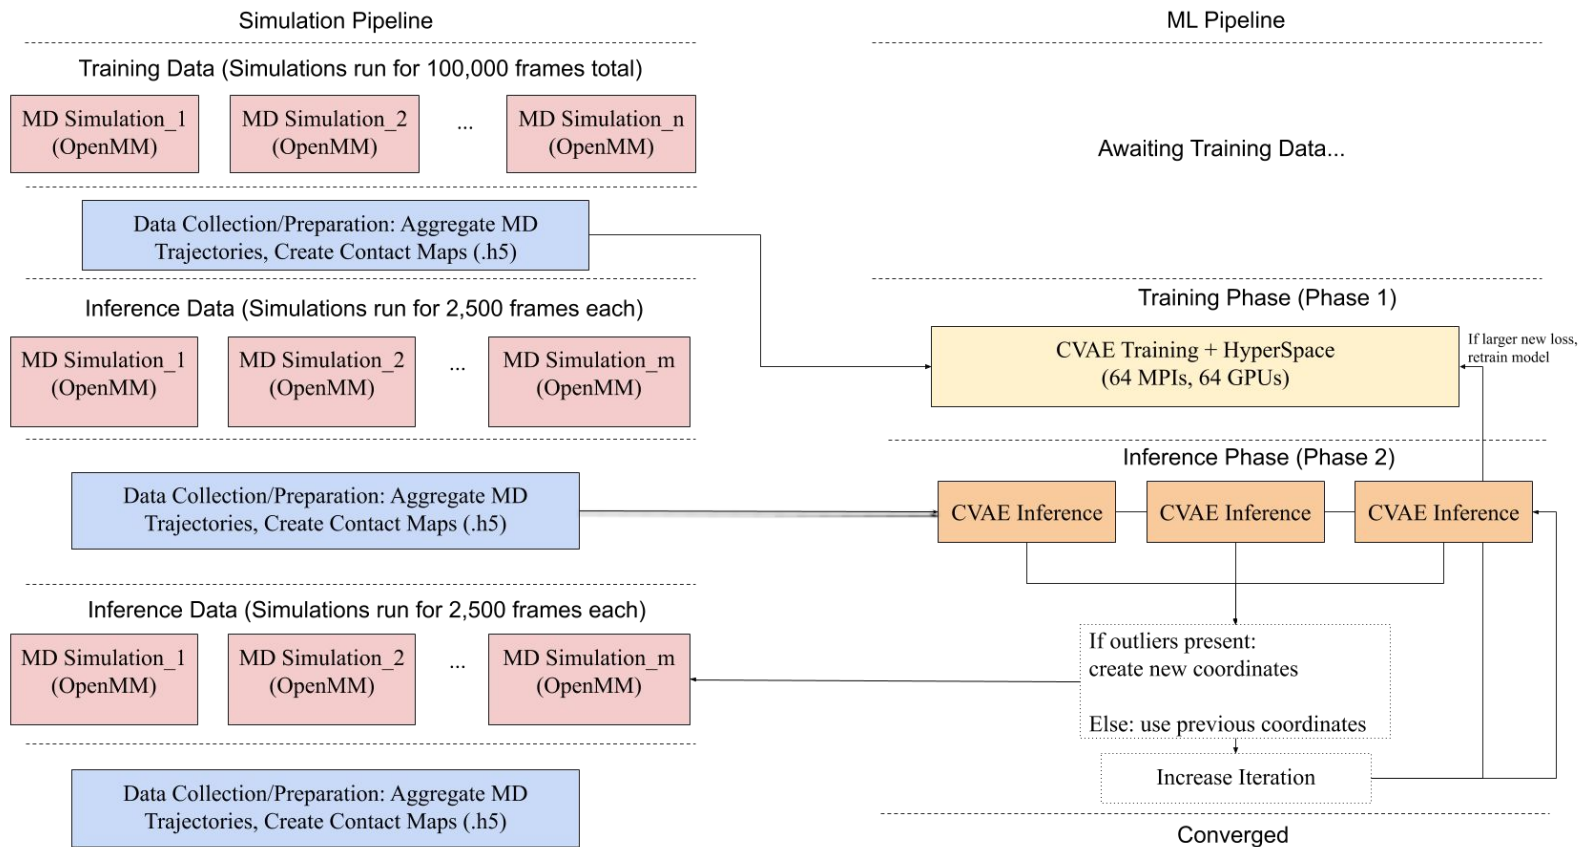
Jan 2019

Matteo Turilli, Andre Merzky, Shantenu Jha

New Use Case: INSPIRE-Candle



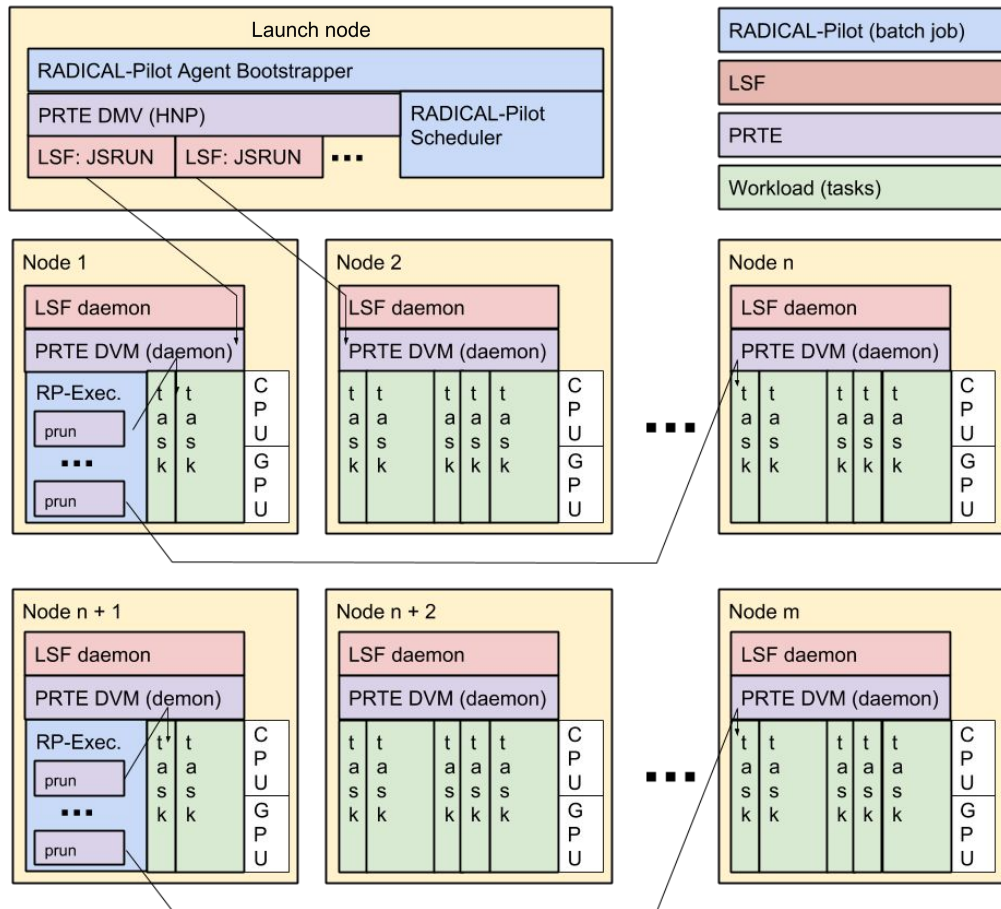
New Use Case: MicroScope



Deploying NGE and RADICAL-Pilot on Summit

- Implemented job submission, task scheduling and task pinning to hardware threads, cores, and GPUs
- Support of concurrent execution of single core, OpenMP, MPI, GPU tasks on the same pilot
- Continuous integration tests for jsrun syntax and for a growing number of RP functionalities
- Issue CCS #401235 is blocking scaling above 1 or 2 nodes, and execution of multiple generations of tasks on the same pilot.
- Alternative to directly using jsrun: ongoing collaboration with ORNL team about using PRRTE on Summit.

Scheduling and Launching Jobs on Summit



- PR RTE project abstracts away LSF and JSRUN
- PRTE DVM
 - Head Node Process
 - Daemons (1 per node)
- PRUN: uses PRTE DVM to place and execute tasks on nodes
- PRTE daemons use PMI-X as interface to LSF (or SLURM, etc.)

Progress of NGE/Harvester Integration Since TIM

- Test runs with bulk task submission by Harvester to NGE pilot
- Performance limitation: Harvester submits ~320 units in about 4000 seconds. This is a rate way below NGE limit.
- Performance limitation: Tasks arrive in small bulks, NGE can manage much larger bulks.
- Test did not produce traces therefore profiling is not available. This may depend on unclean termination or missing/improper configuration.
- Harvester/NGE integration needs dedicated resources on the Harvester side to produce further results.

Planned Activities

- Use cases:
 - Finish ongoing prototype implementation.
 - Run at scale on Summit once CCS #40123 has been addressed.
- Summit support:
 - Coordinate with ORNL and possibly IBM to address issue CCS #40123
 - Continue close collaboration with ORNL to improve PR RTE support
- Harvester connector:
 - Getting resources to support Harvester integration.
 - Increase task submission rate, increase bulk size submission.
 - Perform experiments with relevant workloads.